(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 978 968 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
09.02.2000 Bulletin 2000/06

(51) Int Cl.[7]: **H04L 12/56**, H04Q 11/04

(21) Application number: **99250262.5**

(22) Date of filing: **03.08.1999**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **05.08.1998 US 129662**

(71) Applicant: **Vitesse Semiconductor Corporation**
Camarillo, California 93012 (US)

(72) Inventors:
- **Mullaney, John P.**
  Minneapolis, Minnesota 55436 (US)
- **Lee, Gary M.**
  Camarillo, California 93012 (US)

(74) Representative:
Müller, Wolfram Hubertus, Dipl.-Phys. et al
Patentanwälte
Maikowski & Ninnemann,
Xantener Strasse 10
10707 Berlin (DE)

(54) High speed cross point switch routing circuit with word-synchronous serial back plane

(57) An asynchronous serial crosspoint switch is word synchronized to each of a number of transceiver circuits. The crosspoint switch circuit generates both a master bit clock and a master word clock signal. A transceiver circuit recovers the master bit clock signal from an incoming high-speed serial data stream using a clock and data recovery circuit. The recovered bit clock signal is used as a timing signal by which data is serialized and transmitted to the crosspoint switch circuit. The data stream transmitted to the switch circuit is frequency locked to the master bit clock signal, such that the serial data stream need only be phase adjusted with a data recovery circuit. To recover word timing, the switch circuit issues alignment words to the transceivers during link initialization. The transceivers perform word alignment and establish a local word lock. Alignment words are then reissued to the switch circuit using the local word clock. The switch circuit compares the boundary of the received word clock to the master word clock and, if misaligned, the transceiver shifts its transmitted word by one bit and retries. Necessary edge transition density is provided by overhead bits which also designate special command words asserted between a transceiver and a switch circuit. Flow control information is routed from a receiving transceiver back to the transmitting transceiver using the overhead bits in order to assert a ready-to-receive or a not-ready-to-receive flow control signal. The overhead bits additionally provide information regarding connection requests and other information.

EP 0 978 968 A2

must be made as fast as possible to compensate for transceiver frequency mismatch and to minimize realignment induced dead time. Conventional systems typically use up an additional 10-20% of bandwidth in order to provide a minimum number of transitions to guarantee that the serial data stream comprises a sufficiently high transition rate to support fast phase recovery circuits.

[0009] Serial data transmission may also be synchronous. In synchronous data transmission the sequence of binary "ones" and "zeros", making up the data stream, occurs with reference to a data bit cell defined by a uniform or single-frequency clock signal transmitted with the data. Transmitting the clock signal together with the data, however takes up valuable bandwidth, increases high speed line requirements, and reduces the data transmission capability of the system. In addition, word alignment must still be performed.

[0010] The effects of jitter, or bit shift, in a serial data stream are illustrated in FIG. 2. Data has been phase-locked to a bit clock signal wherein data is stable within a particular bit period such that it may be strobed into an input register on the falling edge of bit clock. Given perfect phase and frequency lock, the periodicity of the bit clock signal might serve to define synchronous bit cells; a logical high data occurring within a code bit cell representing a logic ONE, a logical zero on data occurring within a code bit cell representing a logic zero. The data sequence illustrated would therefore be read as 11011000.

[0011] Phase jitter, frequency mismatch, andior a delay change through the switch matrix, has displaced, or shifted, the serial data stream by approximately 90 degrees in phase. Data stability, of the late data stream, occurs outside of the intended code bit cell, and into the next code bit cell, causing the data stream sequence to be incorrectly read as 01101100 rather than 11011000. Thus, it can be seen that by merely shifting a particular serial data stream by approximately 90 degrees in phase, the binary sequence comprising a data word, as represented within a frame defined by a word clock signal, causes the word to lose all meaning.

[0012] The random nature of data shift can be appreciated by referring to FIG. 3. Shifts in the nominal position of a data transition edge due to timing fluctuations result in a normal distribution of possible transition edges distributed with respect to time around the occurrence of the bit clock timing edge. If the bit clock period is used to define the code bit cell boundaries, there would be an approximately 50% probability that a transition edge, representing a transition from 1 to 0, or 0 to 1, would be shifted early or late and therefore not captured in the proper code bit cell, giving rise to a data word error. A code bit cell should properly have its boundaries symmetric about the mean of an expected data value. However, because of a multiplicity of reference clock signals provided in prior art-type transmission systems, bit cell boundaries must be inordinately wide in order to accommodate the expected transition edge distribution pattern. Widening the bit period necessarily requires that a system bandwidth be consequently reduced, reflecting a loss of transmission capability. Accordingly, some other means must be provided to ensure that all of the component elements of a multiport transmission system be at least frequency locked together, such that only phase recovery is necessary to correctly place the transition edges of a serial data stream within an appropriate code cell boundary.

[0013] The same reasoning holds true for word synchronizing a 2.125 GHz serial data stream. A word detection window (word clock) must be able to accommodate variations in its own frequency and phase in order to provide for accurate detection and capture of a data word from serial data running at slightly variable channel rates. If a word clock signal were to be bit-shifted by the same approximately 90 degrees in phase from the bit clock signal, the same type of data read error would occur as if the bit clock signal were shifted. Thus, it will be understood that in addition to having each of its component elements frequency-locked together, an effective high-speed data transmission system must also provide for a word clock signal which is frequency-locked to a bit clock. Moreover, the word clock signal must accurately define the beginning of a data word and, thus, must be consistent across all of the component elements of such a transmission system.

[0014] In addition, many prior art switches utilize a processing unit to determine a switch configuration and provide flow control signals for controlling the flow of information through the switch. The processing unit receives connection requests and transceiver status signals over a common data bus accessed by all transceivers connected to the switch. Thus, at any given time, transceivers are requesting access to the common data bus to place connection requests and to provide transceiver status signals, such as signals indicating that the transceiver is unable to accept additional data due to the transceivers input buffer being full.

[0015] The use of a common data bus and processor receiving information from a plurality of transceivers over the data bus may result in delays in data communications. That is, the data bus may not be accessible at any given instance due to the data bus already being in use by another transceiver. Accordingly, the system design must take into account delays due to use of a common data bus in determining when to transmit transceiver input buffer full status, and other signals, and additionally switch connections may be delayed due to delays in providing the processing unit the connection requests. Thus, the use of a processing unit and common data bus further decreases the bandwidth of the switch.

SUMMARY OF THE INVENTION

[0016] The present invention therefore provides a method and system for synchronizing data switching

arbitration scheme for allocating switch connection requests. Circuits for implementing round robin arbitration schemes are known in the art. Alternatively, the arbitration logic and switch control circuit could maintain a record of all switching and routing transactions in a port connection table, and thereby identify sender/recipient pairs and keep track of available connections through the switch fabric.

[0021] As will be developed more fully below, routing, or connection, requests (CRQs) are made by a transmitting port to the arbitration logic and switch control circuit 55, which appropriately configures the matrix 53. The logic circuity 56 additionally provides the switch control circuit connection requests overhead bits to effect flow control of data transmitted through the switch.

[0022] A global, system wide clock signal is provided on the switch circuit unit 50 and defines a global word clock (WCLK) signal, which is a 62.5 MHz signal in the described embodiment. A synchronized bit clock timing signal is developed through a CMU circuit 58 using the word clock signal, with the bit clock timing signal being a 2.125 GHz signal in the described embodiment. The WCLK signal is provided by an external 62.5 MHz crystal oscillator, which is coupled to the switch circuit in conventional fashion, but some other suitable reference clock generation circuit may also be used. The bit clock timing signal is directed, globally, to each of the bi-directional switch ports 54 comprising the switch circuit unit 50. Defining the bit clock timing signal for each of the switch ports from a single input reference clock signal (WCLK) has important implications to the synchronous bi-directional data transmission characteristics of the system. Since each of the switch ports operate off of a unitary timing signal developed from a single timing reference, it will be understood that each of the switch ports 54 will operate in a synchronous, albeit possibly phase shifted, fashion with the others.

[0023] Each switch port 54 comprises a receiver section including a serial-to-parallel data converter 60 (also referred to as a deserializer or DMUX). The DMUX is configured to receive a serial data stream transmission from the transmitter section of a corresponding transceiver port card 52 and convert the serial data into a parallel data word (referred to herein as a transmission character). In the embodiment described, the serial-to-parallel converter 60 receives incoming data transmissions at a 2.125 Gb/s data rate and outputs a 34-bit transmission character comprising a 32-bit data word, plus two overhead bits, at a parallel data rate of about 62.5 MHz.

[0024] Similarly, each switch port 54 comprises a transmitter section including parallel-to-serial data converter 62 (also referred to as a serializer or MUX) which performs a similar function to the serializer 60, but in reverse. The parallel-to-serial converter 62 receives a 34-bit transmission character (a 32-bit data word, plus two overhead bits) which has been routed to the corresponding switch port through the switch fabric 53 at an input parallel data rate of approximately 62.5 MHz. The parallel-to-serial converter converts the parallel data into a serial data stream suitable for transmission to a receiver portion of the port card 52 at a serial data rate of approximately 2.125 Gb/s. Thus, the 62.5 MHz WCLK signal, distributed by the CMU clock circuit 58, is used as a master strobe to clock 34-bit parallel data out of the serial-to-parallel converter 60 to the port logic circuit 56. The WCLK signal is also used to clock 34-bit parallel data from the port logic circuit 56 into the parallel to serial converter 62. All timing signals, whether serial bit timing signals or parallel word timing signals, used by each of the switch ports, therefore, is developed by the CMU clock circuit 58 in response to the system wide reference WCLK.

[0025] Each of the transceiver port cards 52 are typically constructed to include a mix or combination of transceiver circuitry and circuitry related to a particular user's application. In a typical configuration, a transceiver port card includes physical layer circuitry 61, 63 for a given communication protocol, and data buffer circuitry that manages the information flow and formatting between downstream user application circuitry and the transceiver. In the embodiment of FIG. 4, the data buffer circuitry comprises transmit and receive FIFOs 64 and 66, respectively. The transmit and receive FIFOs are each coupled to the transceiver circuitry over a parallel data interface. In the embodiment described, data is clocked to the transceiver at the 62.5 Mb/s parallel data rate.

[0026] In many communications applications, however, the parallel data interface to the transmit and receive FIFOs will operate at a different frequency than the 62.5 MHz word clock used by the switch card and the transceiver port card's transmit and receive circuitry. In this case, the transmit and receive FIFOs 64 and 66 are implemented as synchronous, dual port FIFOs, whose dual clock ports are used to elastically span any discontinuous clock boundaries between the transmission side and the media side. In addition, the transmit and receive FIFOs are made large enough to function as data queues or data buffers for each port card's transceiver circuitry. The transmit and receive FIFOs may be implemented as register stacks, string buffers, and the like, but are preferably implemented as dual-port, parallel data buffer, integrated circuit memory elements. Suitable transmit and receive FIFOs 64 and 66 include FIFO devices able to operate at speeds up to 67 MHz. Such a FIFO is exemplified by the IDT7236 series of synchronous FIFOs, manufactured and sold by Integrated Device Technology, among others.

[0027] Each transceiver port circuity 68 comprises a transmitter section, indicated as TX, and a receiver section, indicated as RX. The transmitter section includes a serializer (not shown in the embodiment of FIG. 4) for converting 62.5 Mb/s parallel data from the transmit FIFO 64 into a serial data stream suitable for transmission to the receiver section of a corresponding switch

port and thence to its corresponding intended recipient transceiver port. This is done under the operational control of an arbitration logic and switch control circuit 55. As referred to previously in the embodiment of FIG. 4, the switch matrix or fabric 53 is conventional in implementation and design, and need not be further described herein. It should be noted, however, that unlike conventional switches implementing a switch fabric, the arbitration logic and switch control circuitry 55 is not implemented as a conventional central control processor. However, a control processor may be used with various aspects of the present invention.

[0035] Further, flow control decisions are not delayed by using data provided by the transceivers and routed to a conventional central control processor using a dedicated data bus. Instead, overhead bits are sent through essentially a reverse crosspoint switch and appended to the 32-bit data packet (thereby defining a 34-bit transmission character) and used to directly control the flow of information from a transmitter to a receiver, as well as provide other information.

[0036] Turning now to the transceiver port circuitry 68 of FIG. 5a, the transceiver can be viewed as comprising two parts, a transmitter section and a receiver section. When the transceiver is in transmit mode, a 32-bit data word TXIN[31:0] and a 2-bit transmission type word TXTYP[1:0], are clocked out of the transmit FIFO 64 and into a synchronizing parallel input buffer register 70 in accord with 62.5 MHz clock rate.

[0037] Parallel data is clocked out of the buffer register 70 and into a transmit control logic circuit 71. Transmit control logic circuit 71 is responsible for asserting transmission state signals to the transmit FIFO. Such state signals include indications that data packets have been successfully transmitted to all outputs, a beginning-of-packet indication, a retransmission required indication, and the like. Also, transmit control logic 71 is responsible for adaptively reconfiguring TXTYP[1:0] information into a 2-bit flow control overhead bit field when the transceiver is configured to operate in a particular communication mode, designated "overhead-mode" herein. An overhead-mode signal, OH-MODE, is a user programmable state signal, externally sourced, and coupled to the transmit control logic circuitry 71 over an internal communication bus 75.

[0038] In any communication mode, the transmit control logic 71 combines the 32-bit data TXIN[31:0] with either the 2-bit TXTTYP[1:0] or a 2-bit flow control overhead bit field, into a 34-bit data string, the 34-bit data string comprises a transmission word or transmission character. The 34-bit wide transmission characters are serialized by a parallel-to-serial converter 72 (also referred to as a serializer or MUX) and provided to a 2.125 GHz serial output buffer 74. Serialized data is clocked out of the transceiver 68 as a differential signal TXS+/TXS- over a high speed serial transmission line to the input of a corresponding switch port 54. The serial output buffer 74 is clocked by a 2.125 GHz bit clock signal which is, in turn, directly developed by the CRU from an incoming serial data stream sent by the switch port 54. The bit clock signal, BCLK, is a 2.125 GHz strobe which defines the bit cell boundaries of the desired serial data stream. The BCLK signal is directed through a timing generator 86 which comprises divide-by-34 circuitry, such that the 2.125 GHz BCLK signal is divided down to the 62.5 MHz transceiver transmit word clock signal, which is also the transceiver master word clock WCLK signal, in synchronous fashion. Thus, it will be understood that the word boundaries of WCLK and, thus, each 34-bit wide transmission character, will correspond to and be synchronous with every 34th strobe transition edge of the BCLK signal.

[0039] On the receiver side, a high speed serial transmission line is coupled between a high speed output of the switch port 54 and the receiver input of the transceiver circuit 68. The transmission line is configured to provide a differential, serial data stream RXS+/RXS- to the transceiver 68 at a 2.125 Gb/s data rate. The receiver input is coupled to a deserializer, or serial-to-parallel converter 78 which suitably converts the 2.125 GHz serial data stream into 62.5 MHz 34-bit wide parallel transmission characters.

[0040] Serial data is transmitted by the switch port 54 for retrieval by the transceiver circuit 68 without any additional timing reference signals added thereto. A serial stream of data flows over the transmission line with no accompanying clock information. However, the deserializer 78 must process the serial data stream synchronously, such that the resulting 34-bit wide parallel transmission characters are correctly aligned on the appropriate word boundaries. Thus, timing information, i.e., a clock signal, is recovered directly from the serial data stream by a clock recovery unit (CRU) 80. The CRU 80 is a phase and frequency sensitive clock recovery circuit, such as a high-speed phase locked loop (PLL). PLL circuitry suitable for extracting a 2.125 GHz BCLK signal from a 2.125 Gb/s data stream are common circuits implemented in high speed transceiver applications and are well understood by those having skill in the art. Accordingly, it is considered unnecessary to go into detail regarding their construction and operation herein. It is sufficient that CRU 80 is able to recover a bit clock signal BCLK from a serial data stream provided by the switch port 54, and that the recovered BCLK signal is frequency-locked to the frequency of the serial data stream transmitted by the switch port.

[0041] The recovered clock signal, BCLK, is directed through a divide-by-34 timing generator 82 which provides a 62.5 MHz transceiver receive word clock signal to the deserializer 78 and input side of retiming bank 79. The output side of the retiming bank receives the transceiver master word clock signal WCLK, as do receive control logic circuitry 83, a parallel output buffer register 84 and a receiver FIFO 66, from whence received transmission characters are directed to follow-on customer application circuitry over a parallel data interconnect

invention, it is nevertheless necessary to also provide for some means to word synchronize the information communicated between transceiver 52 and the switch port 54. Even though the transceiver 52 and switch port 54 are frequency-locked together, the transmit and receive data streams may be out of word alignment, with resultant loss of transmission character content. Accordingly, word (or frame) alignment must be established and maintained throughout serial data transmission. To recover word timing, the switch circuit 50 issues particular, pre-defined alignment words to each of the transceivers 68 during transmission link initialization and handshake protocol establishment.

[0047] Referring now to FIGS. 5a and 5b, and the flow charts of a word alignment and synchronization process illustrated in FIG. 6a, word timing synchronization is established between a switch port and its associated transceiver port by an adaptive feed-back process. In the adaptive feed-back process a predefined alignment word transmitted by the switch are used by the transceivers to establish a transmitter receive word clock. The transceiver then transmits alignment words to the switch, with the switch then comparing the alignment word to an expected alignment word. The switch continues to issue alignment words in the event that the alignment word does not match with the transceiver shifting its transmitted alignment word in bit-by-bit fashion until the transceiver is word synchronized to the switch.

[0048] A flow chart of the word alignment process is illustrated in FIG. 6A. The word alignment process occurs upon power up, reset, or link initialization. The word alignment process executes independently for each transceiver. The word alignment process, therefore, may execute in parallel for any number of transceivers. In step 100 of the word alignment process,the transceiver transmits at least one reset word to the switch. The reset word, comprising all logic "ones" in the described embodiment, requests that the switch begin the initialization and word synchronization process. The receipt of a reset word by the switch causes the switch to transmit alignment words to the transceiver. Alignment words are generated in Step 120 by an alignment word generator and comparator 100 comprising a portion of the port logic circuitry 56. In the embodiment presently described, upon receipt of the reset word, the port logic circuitry 90 causes the alignment word generator and comparator circuit 100 to sequentially generally output alignment words through the serializer 92 and serial output buffer 94. The alignment words are transmitted to the receiver input of the corresponding transceiver port circuitry 68. Alignment words include no inherent data content and so may be devised to contain any form of binary information. Preferably, alignment words are encoded such that the word (frame) boundaries can be easily determined by the alignment word generator and comparator circuit 100. Such an encoding scheme may be implemented in an alignment word comprising a "1" in the LSB and MSB positions of the word, with the remaining bit cells comprising a "0" string, i.e., 10000001, using an 8-bit word as an example. Other bit patterns with increased edge density, such as 10101011, may also be used.

[0049] In step.101 the process determines if the transceiver detects the alignment word. The transceiver circuitry which accomplishes this deserializes alignment words using the deserializer 78 and uses an alignment detector circuit 102, coupled to "snoop" the parallel bus coupled between the deserializer 78 and the transceiver's retiming register bank.

[0050] If the alignment detector does not detect the correct alignment word the alignment detector provides a signal to the receive word clock timing generator to shift the receive word clock by one bit in step 102. Examination of the received alignment words and, if necessary, the shifting of the receive word clock continues until the alignment detector detects the correct alignment word. Once the receive word clock is correctly aligned, the transceiver begins transmitting alignment words to the switch in step 103. The switch, using the alignment word generator and comparator circuit 100 compares the received alignment words to the expected alignment word in step 104. If the switch detects the correct alignment word the switch in step 107 issues IDLE word to the transceiver to signal that the transmitter is now word synchronized with the switch. The process then returns.

[0051] If the transceiver continues to receive alignment words, the alignment detection circuit 102 causes the frame generator and bit shifter to shift its transmit word boundary by one bit position after receipt of every 32 alignment words from the switch in steps 105 and 106. The process repeats until the alignment word generator and comparator circuit 100 in the switch determines that the alignment words sent by the transceiver are correctly framed in accordance with the switch word clock signal.

[0052] Bit and word alignment is accomplished only during link initialization. Since this process occurs relatively infrequently, the system of the present invention is not required to support fast phase acquisition and need only maintain frequency lock in the manner described above. Moreover, any small variation in phase of the signal received by the switch, whether due to component aging or temperature variations, is accounted for by the data recovery unit (DRU) of the switch circuitry. In addition, since a master reference clock is provided on the switch circuit, the system does not require transmission characters to incorporate additional overhead bits devised to absorb the types of bit loss that can occur with multiple reference clocks driving multiple transceivers, as is common in prior art implementations.

[0053] Although ensuring that the switch and all of its attendant transceiver circuits are bit and word synchronized, the synchronization method does require some means to ensure that both transmit and receive serial

in a 16-bit active connection field C[15:0].

[0060] Thus, in accordance with the invention, a command word can be used to send a connection request (CRQ) from a particular transceiver port to the switch matrix. An acknowledgment (ACK) to the request is returned from the switch to the requesting transceiver port by appending two overhead bits, configured (1,1) in either a command word's overhead field 114 or a data word's overhead field 112. Flow control channel information is also shared between a receiving transceiver port and a transmitting transceiver port by reconfiguring the two overhead bits comprising the overhead bit fields 112 and 114 of a data or command word. As will be described in greater detail below, a ready-to-received (RTR) configuration signal informs the transmitting transceiver port that there is sufficient room in the receiving transceiver port's receive FIFO and that it is appropriate to continue to transmit data. A not-ready-to-receive (NRTR) signal informs the transmitting transceiver port that the receiving transceiver port's receive FIFO is filling up and it is, therefore, not appropriate to continue sending data.

[0061] A ready-to-receive (RTR) signal is generated by configuring the overhead bits as (0,1) while a not-ready-to-receive (NRTR) signal is generating by configuring the overhead bits as (1,0). When the overhead bits are configured as (1,1), as indicated above, the overhead bits comprise an acknowledge (ACK) signal which is returned from the switch to a transceiver port which has made a connection request. It is the function of the switch, particularly the port logic 90 and arbitration logic and switch control circuit 55 to either generate the appropriate overhead bits (such as ACK) or to intercept flow control messages (such as RTR and NRTR) and re-direct them to the appropriate transceiver port so that effective flow control is maintained.

[0062] In order to better understand the utility of the overhead bits, it will be useful to consider how flow control overhead bits are generated by a transceiver in response to a FIFO almost full condition, and how flow control overhead bits are used in order to prevent further transmission. Referring now to FIGS. 5a and 5b, the receive FIFO 66 is conventionally provided with a signal line that indicates that the FIFO is in an almost full condition. The receive FIFO 66 can overflow if data arrives and is written to the FIFO faster than it can be read from the FIFO to the user's application circuit or physical media. In the example of FIG. 5a, the almost full signal AF from the receive FIFO 66 is coupled to the transceiver's input register 70 which passes the AF signal to the transceiver's transmit control circuit 71. The AF signal will be asserted when the number of empty FIFO locations is less than or equal to a value pre-programmed into the FIFO. This minimum value typically depends on the latency period between the time when the AF signal is asserted and when it is received at the transmitting transceiver. The latency period is such that 14 + N more words are able to be transmitted, where N relates to a distance latency parameter and depends on the physical distance between the transceivers and the switch. It will be understood by those having skill in the art how to calculate the number of words that might be written into an almost full FIFO during a latency period and how to program this value into the FIFO such that AF is asserted at the proper time.

[0063] Transmit control circuit 71 receives the almost full indication and, if put into overhead mode by the appropriate OH-MODE signal, appends a not-ready-to-receive (NRTR) signal to a command or data word. As mentioned above, an NRTR signal is generated by configuring the overhead bits as (1,0). The command or data word including the NRTR signal is transmitted to the switch which recognizes the overhead bits as not comprising a conventional pattern and, in response, strips the NRTR signal from the command or data word and routes it to the appropriate transceiver for action. The overhead bits are recognized by the switch port's port logic circuit 56 which directs them to the switch's arbitration logic and switch control circuit 55 through its parallel bus connection through a plurality of 16 to 1 MUXes 91 (only one of which is shown for clarity). The 16 to 1 MUXes provide essentially the functions of a reverse cross-point switch, which is how the described function is implemented in one embodiment. Arbitration logic and switch control circuit 55 recognizes which of the 16 switch ports provided the NRTR signal and, since it is in control of the configuration of the fabric 53, the arbitration logic and switch control circuit 55 understands which of the 16 switch ports is coupled to the transceiver which is transmitting data to the almost full recipient. Arbitration logic and switch control circuit 55 then provides the NRTR signal (1,0) to the appropriate switch port's port logic circuit 56 through MUX 91. That switch port's port logic circuit 56 appends the NRTR overhead bits to the next outgoing transmission to its corresponding transceiver 68.

[0064] In the transceiver port circuitry 68, the (1,0) overhead bits are directed to the receive control circuitry 83, as described above, which recognizes that the (1,0) overhead bit pattern represents a not-to-ready to receive condition and that the transceiver should cease transmitting. Receive control circuit 83 communicates with transmit control logic circuit 71 over the internal communication bus 75 and passes the overhead signal values to the transmit control logic circuit for command processing. In transmit control circuit, the received flow control signal is ANDed with a read enable signal, controlled by the transmit control logic. The resulting read enable signal REN is connected through the input register 70 to the transmit FIFO 64. When a not-to-ready to receive signal is ANDed with read enable, the resulting REN signal is de-asserted, instructing the transmit FIFO 64 that data reading is no longer enabled. Data transmission thereby ceases. Thus, the transceiver is only able to read a word from the transmit FIFO 64 if the flow control signal (the overhead bits) from the receiving

to the at least one switch port for detecting a received alignment word received by the at least one switch port.

2. The high speed communication switch circuit of claim 1 wherein the switch ports comprise receive switch ports for receiving data and transmit switch ports for transmitting data.

3. The high speed communication switch circuit of claim 2 further comprising plurality of alignment word detectors, each alignment word detector operatively coupled to a corresponding receive switch port.

4. The high speed communication switch circuit of claim 3 further comprising a plurality of data recovery units operatively coupled to a corresponding receive switch port for recovering data received by the receive switch port.

5. The high speed communication switch circuit of claim 4 further comprising an arbitration circuit for configuring the switch fabric and port logic circuitry associated with each receive switch port for routing connection requests to the arbitration circuit.

6. The high speed communication switch circuit of claim 5 wherein the data comprises payload and overhead bits.

7. The high speed communication switch circuit of claim 6 further comprising routing means for routing the overhead bits received by a receive switch port to a transmit switch port, and wherein the port logic circuitry includes means for detecting the overhead bits received by the receive switch port and means for providing the routing means the overhead bits.

8. A method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system comprising:

providing a reference clock signal from a clock signal generator;
forming a word cell boundary signal using the reference clock signal;
transmitting a predefined transmit data word to a receiver, the data words having bit periods defined by the reference clock signal and word cell boundaries defined by the word cell boundary signal;
receiving a receive data word from a receiver; and
determining if the receive data word matches a predefined data word.

9. The method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system of claim 8 further comprising continuing to transmit the predefined data word to the receiver until receiving a receive data word from the receiver which matches the predefined data word.

10. The method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system of claim 9 further comprising transmitting an idle word when the receive data word matches the predefined data word.

11. A method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system comprising:

transmitting a reset request signal;
receiving a signal containing repetitions of a first predefined alignment word;
extracting a clock signal from the signal containing repetitions of the first predefined alignment word;
generating a receive word cell boundary clock signal based on the clock signal;
extracting a test alignment word from the signal containing repetitions of the first predefined alignment word using the clock signal and the receive word cell boundary clock signal; and
determining if the test alignment word matches the first predefined alignment word.

12. The method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system of claim 11 further comprising shifting the word cell boundary clock signal by a predefined period of the clock signal if the test alignment word does not match the predefined alignment word.

13. The method of synchronizing clock signals of transceivers and switches in a high speed data communication switching system of claim 12 further comprising:

generating a transmit word cell boundary based on the clock signal;
transmitting a second predefined alignment word using the transmit word cell boundary;
determining if a predefined idle word has been received after a predefined time period subsequent to transmitting the second predefined alignment word; and
shifting the transmit word cell boundary by a predefined period of the clock signal if no predefined idle word has been received after the predefined time period subsequent to transmitting the second predefined alignment word.

25. The high speed communications transceiver circuit of claim 24 further comprising an alignment word detector for comparing a contiguous number of bits received by the serial input port with a predefined bit pattern.

26. The high speed communications transceiver circuit of claim 25 further comprising a receive word timing generator operatively coupled to the alignment word detector and the clock recovery unit for defining a receive word cell boundary for contiguous bits received by the serial input port.

27. The high speed communications transceiver circuit of claim 26 further comprising a transmit word timing generator operatively coupled to the alignment word detector and the clock recovery unit for defining a transmit word cell boundary for contiguous bits transmitted by the serial output port.

28. A switching system comprising:

a switch circuit including a plurality of switch ports transmitting and receiving data, a switch fabric coupled to the plurality of switch ports, the switch fabric providing transmission channels for routing data among and between the switch ports, and a timing reference signal defining bit and word cell boundaries for the data transmitted by each switch port; and
a plurality of transceiver circuits linked to a corresponding one of the plurality of switch ports, the transceiver circuits receiving data from and transmitting data to the switch ports, each transceiver circuit comprising:

means for recovering a timing signal from data received from the switch port, the timing signal representing bit cell boundaries of the data as defined by the switch circuit;
means for recovering word cell boundaries from a data received from the switch port, such that each transceiver circuit's word cell boundaries correspond to the switch circuit word cell boundaries defined by the timing reference signal; and
means for transmitting data to the switch port, the timing signal defining bit and word cell boundaries for the data transmitted by the transceiver circuit.

29. A high speed self-routing network switching system in which internal timing reference signals of a plurality of transceiver circuits are word synchronized to a master word clock timing reference signal developed by a switch circuit, the switch circuit including a plurality of switch ports each coupled to communicate with a corresponding transceiver circuit, and a switch fabric for routing data among and between the switch ports, the system comprising:

a first alignment word generator circuit disposed in the switch, the first alignment word generator defining alignment words for transmission to a transceiver, each alignment word generated in accord with word boundaries defined by the master word clock timing reference;
an alignment word detector, disposed in each transceiver, the alignment word detector adjusting a receive word clock boundary until the alignment word detector detects proper alignment, the alignment word detector thereby defining a local word clock timing signal; and
a second alignment word generator circuit, disposed in each transceiver circuit, the second alignment word generator defining alignment words for transmission to the switch in accord with word boundaries defined by the local word clock timing signal, the switch detecting alignment words received from a transceiver with the master word clock timing reference.

30. In a high speed network switching system in which internal timing reference signals of a plurality of transceiver circuits are frequency locked to a timing reference signal developed by a switch circuit, the switch circuit including a plurality of switch ports each coupled to communicate with a corresponding transceiver circuit, and a switch fabric for routing data among and between the switch ports, a method for word synchronizing a transceiver to a master word clock timing reference, the method comprising:

generating alignment words in the switch for transmission to a transceiver, each alignment word generated in accord with word boundaries defined by the master word clock timing reference;
receiving alignment words by an alignment word detector disposed in a transceiver;
adjusting a receive word clock boundary until the alignment word detector detects proper alignment of the alignment words; and
defining a receive word clock timing signal when the alignment word detector detects proper alignment of the alignment words;

31. The method of claim 30 further comprising:

defining alignment words in the transceiver for transmission to the switch in accord with word boundaries defined by a transmit word clock timing signal;
detecting alignment words received from a

ers adapted to transmit serial data formed in data words, the data words comprising overhead bits and payload, to the switch receive ports and to receive serial data formed in data words from the switch transmit ports, the switch including:

a switch fabric reconfigurably interconnecting the switch transmit and receive ports such that the switch fabric provides a switch transmit port of a first pair of switch transmit and receive ports with payload from a switch receive port of a second pair of switch transmit and receive ports and the switch fabric provides the switch transmit port of a third pair of switch transmit and receive ports with payload from switch receive port of the first pair of switch transmit and receive ports;

a reverse switch fabric reconfigurably interconnecting the switch transmit and receive ports such that the switch fabric provides the switch transmit port of the first pair of switch transmit and receive ports with overhead bits from the switch receive port of the third pair of switch transmit and receive ports and the switch fabric provides the switch transmit port of the second pair of switch transmit and receive ports with overhead bits from the switch receive port of the first pair of switch transmit and receive ports; and

a means for routing the payload to the switch fabric and means for routing the overhead bits to the reverse switch fabric.

36. The high speed network switching system of claim 35 wherein the overhead bits transmitted by a particular transceiver indicate that the particular transceiver is able to receive data when the overhead bits are configured in a first pattern.

37. The high speed network switching system of claim 35 wherein the overhead bits transmitted by a particular switch transmit port to a particular transceiver indicate acknowledgment of receipt by the particular switch receive port of a request by the particular transceiver.

38. The high speed network switching system of claim 35 wherein the overhead bits transmitted by a particular switch transmit port to a particular transceiver indicate the granting of a connection request requested by the particular transceiver.

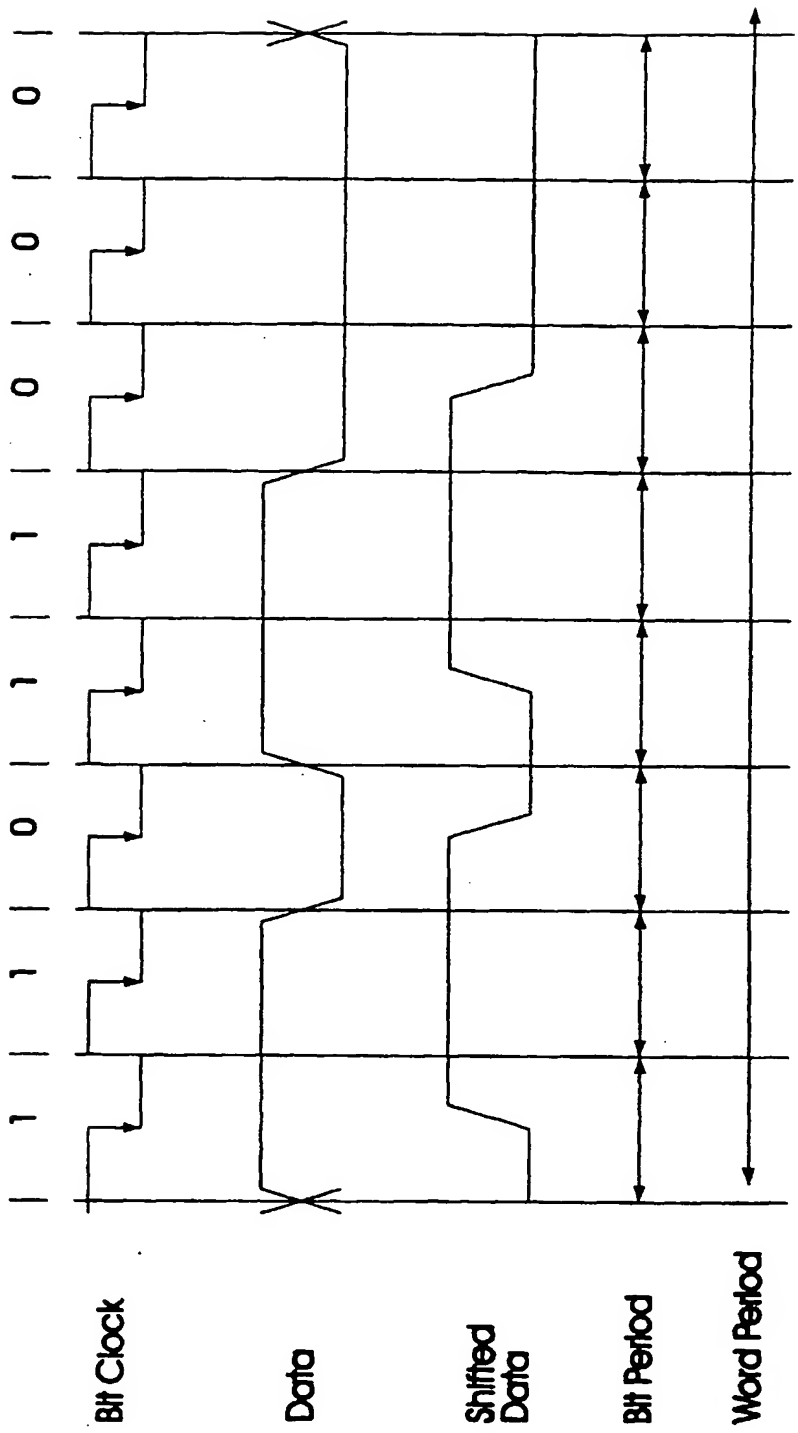

FIG. 2

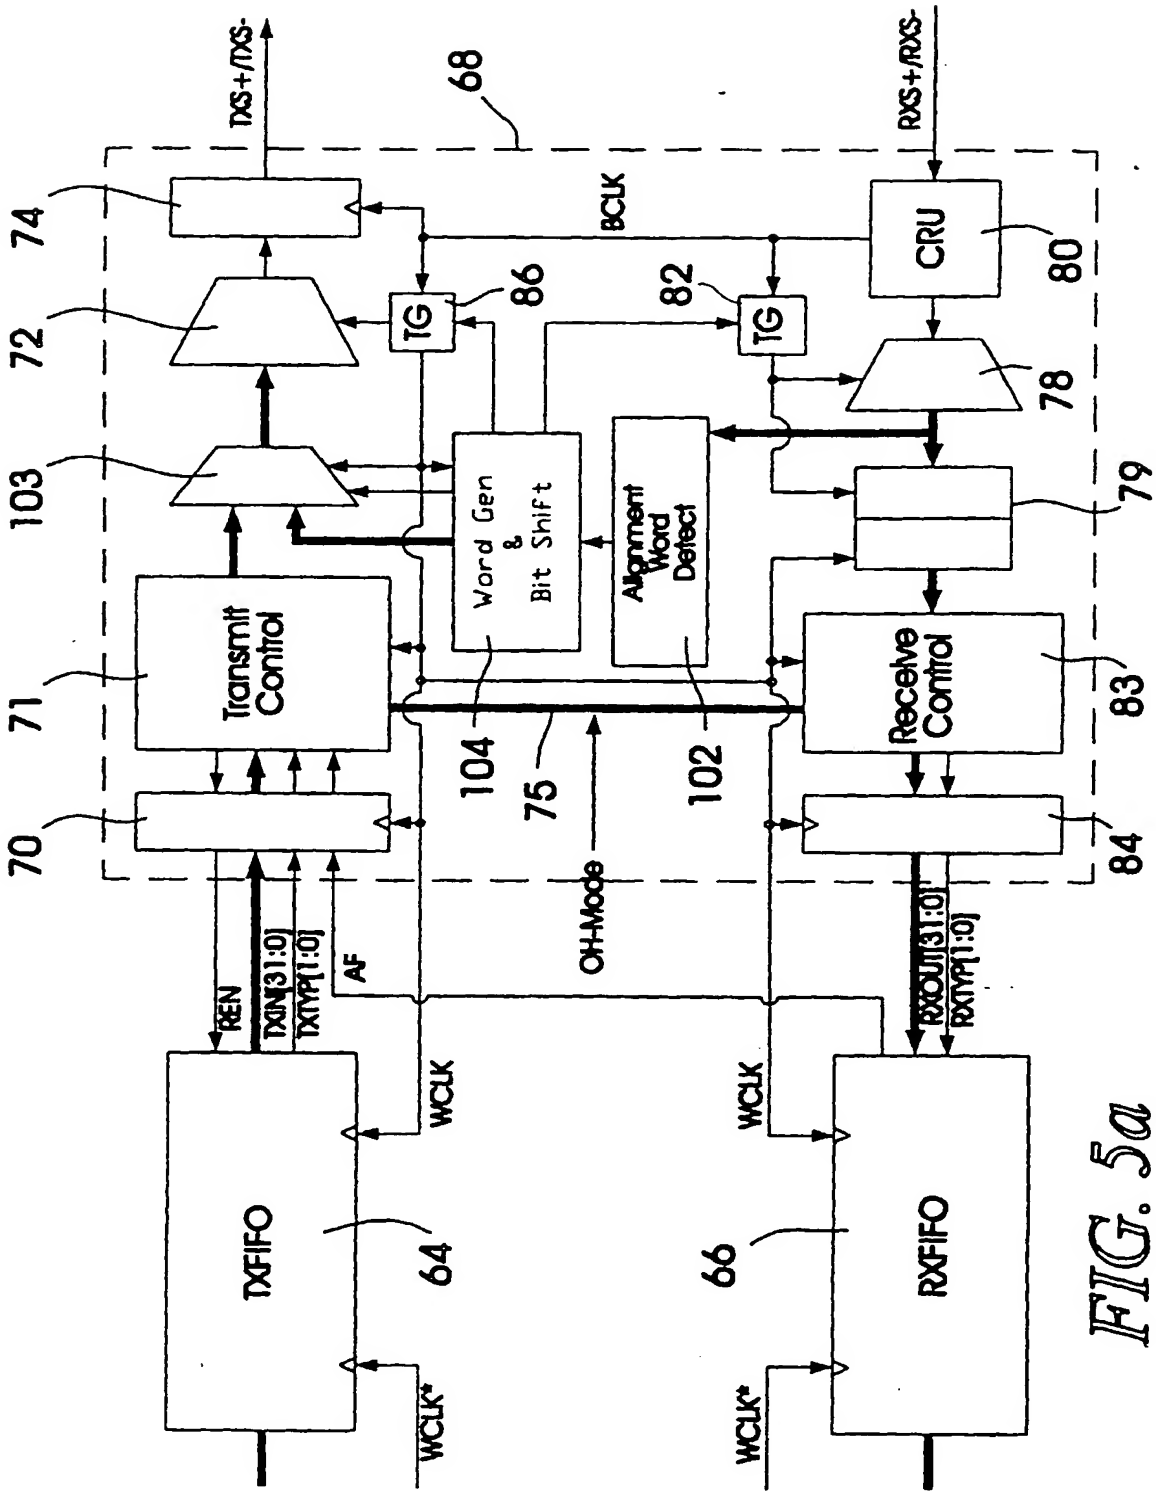

FIG. 5a

FIG. 6

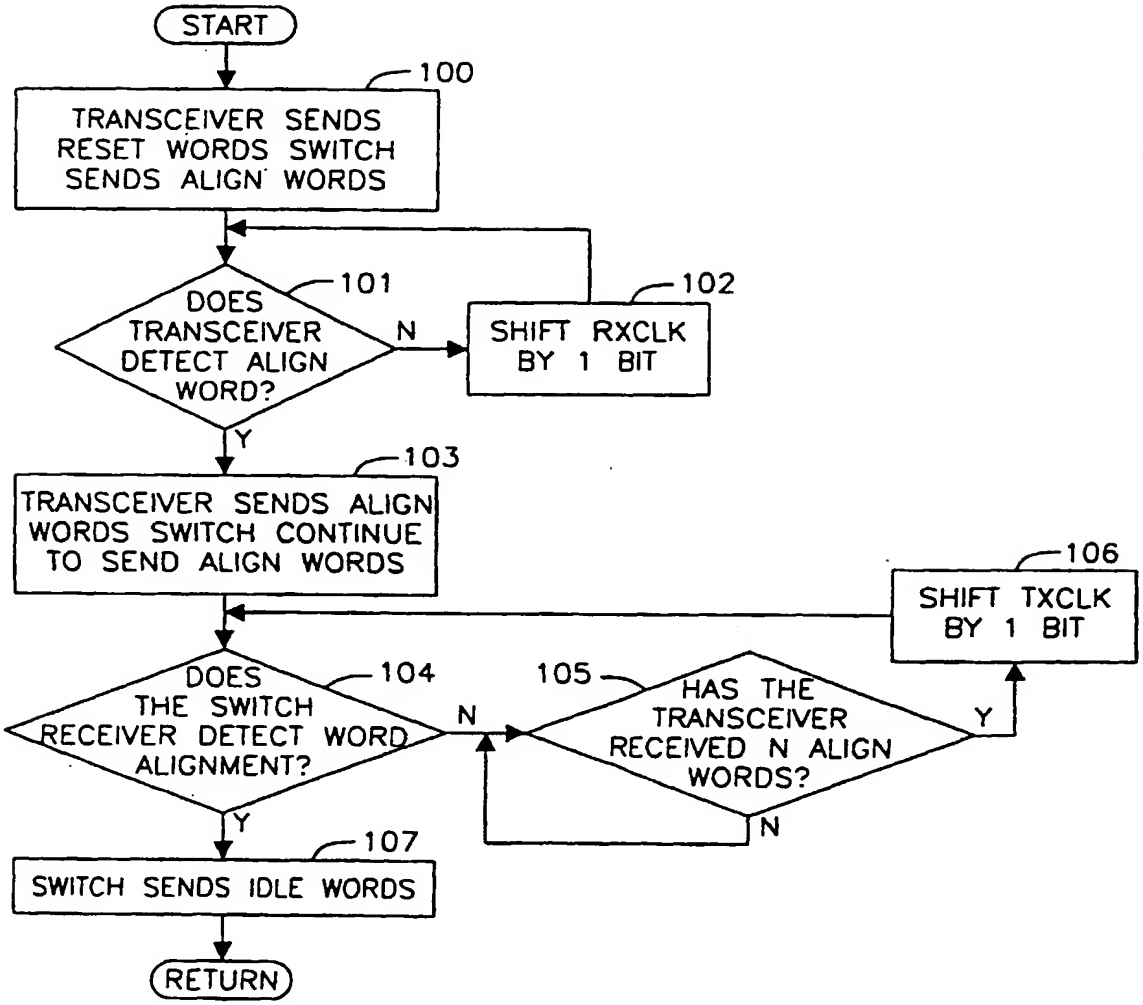

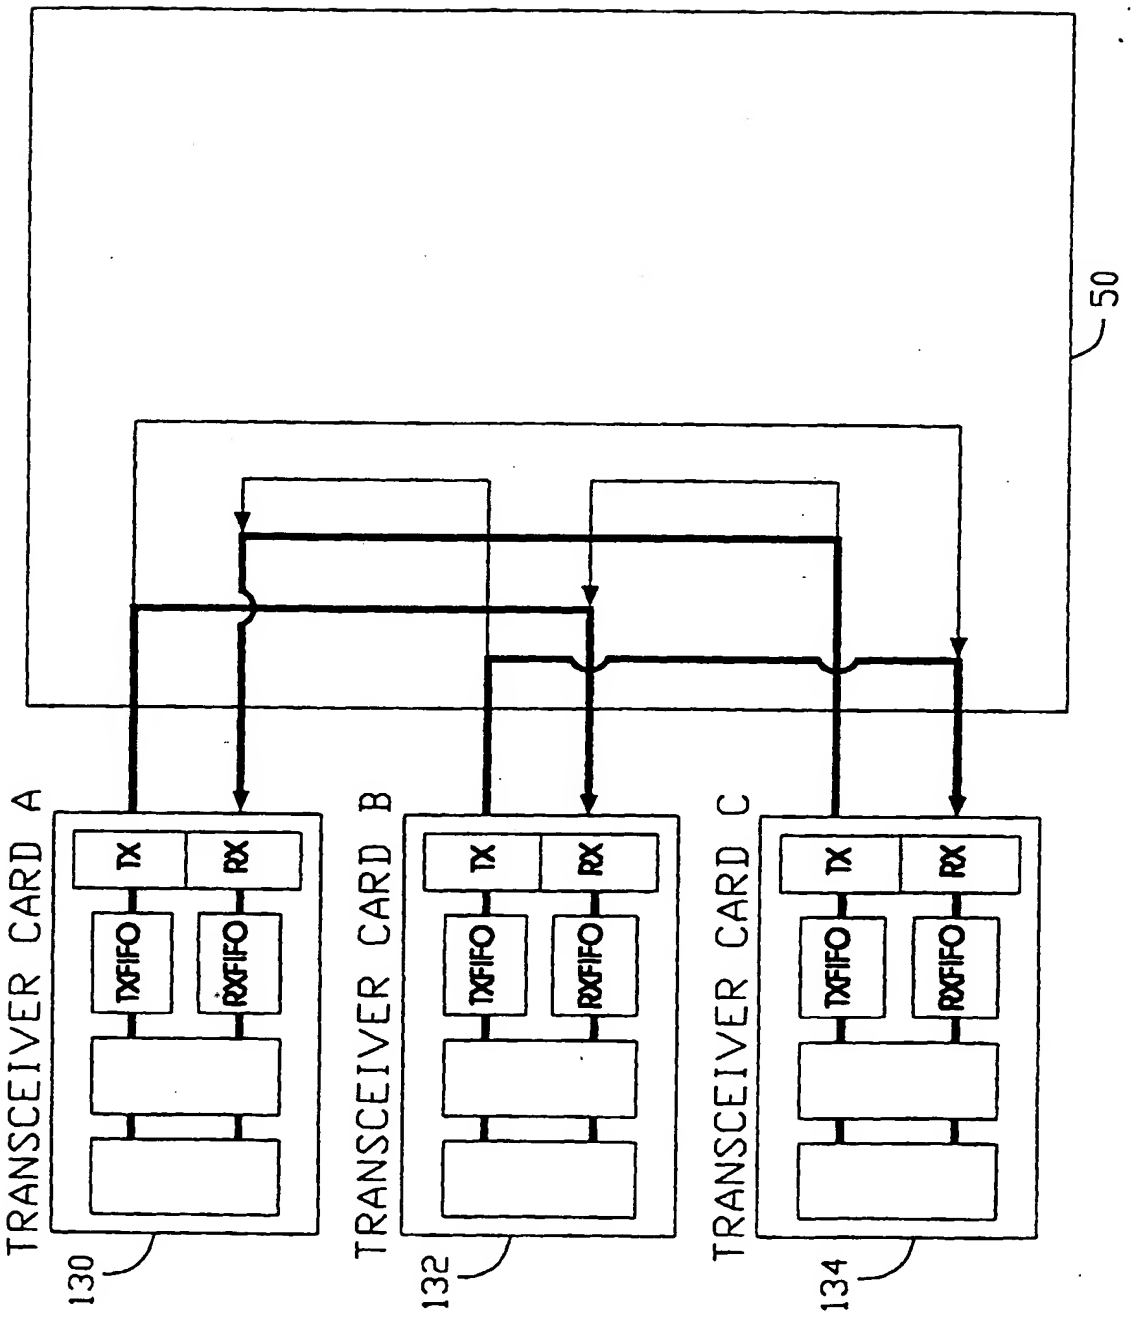

FIG. 8

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 978 968 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
23.04.2003 Bulletin 2003/17

(43) Date of publication A2:
09.02.2000 Bulletin 2000/06

(21) Application number: **99250262.5**

(22) Date of filing: **03.08.1999**

(51) Int Cl.[7]: **H04L 12/56**, H04Q 11/04

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **05.08.1998 US 129662**

(71) Applicant: **Vitesse Semiconductor Corporation**
Camarillo, California 93012 (US)

(72) Inventors:
- **Mullaney, John P.**
  Minneapolis, Minnesota 55436 (US)
- **Lee, Gary M.**
  Camarillo, California 93012 (US)

(74) Representative:
Müller, Wolfram Hubertus, Dipl.-Phys. et al
Patentanwälte
Maikowski & Ninnemann,
Postfach 15 09 20
10671 Berlin (DE)

(54) High speed cross point switch routing circuit with word-synchronous serial back plane

(57) An asynchronous serial crosspoint switch is word synchronized to each of a number of transceiver circuits. The crosspoint switch circuit generates both a master bit clock and a master word clock signal. A transceiver circuit recovers the master bit clock signal from an incoming high-speed serial data stream using a clock and data recovery circuit. The recovered bit clock signal is used as a timing signal by which data is serialized and transmitted to the crosspoint switch circuit. The data stream transmitted to the switch circuit is frequency locked to the master bit clock signal, such that the serial data stream need only be phase adjusted with a data recovery circuit. To recover word timing, the switch circuit issues alignment words to the transceivers during link initialization. The transceivers perform word alignment and establish a local word lock. Alignment words are then reissued to the switch circuit using the local word clock. The switch circuit compares the boundary of the received word clock to the master word clock and, if misaligned, the transceiver shifts its transmitted word by one bit and retries. Necessary edge transition density is provided by overhead bits which also designate special command words asserted between a transceiver and a switch circuit. Flow control information is routed from a receiving transceiver back to the transmitting transceiver using the overhead bits in order to assert a ready-to-receive or a not-ready-to-receive flow control signal. The overhead bits additionally provide information regarding connection requests and other information.

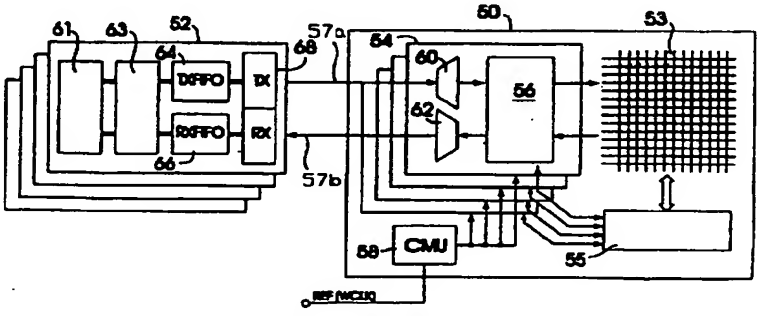

FIG. 4

EP 0 978 968 A3

European Patent Office

Application Number
EP 99 25 0262

CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

☐ Only part of the claims have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid, namely claim(s):

☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

see sheet B

☒ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.

☐ As all searchable claims could be searched without effort justifying an additional fee, the Search Division did not invite payment of any additional fee.

☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid, namely claims:

☐ None of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims, namely claims:

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 99 25 0262

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

25-02-2003

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
WO 9642158	A	27-12-1996	SE	506540 C2	12-01-1998
			AU	6021896 A	09-01-1997
			BR	9608762 A	06-07-1999
			CA	2224196 A1	27-12-1996
			CN	1192836 A ,B	09-09-1998
			EP	0872085 A1	21-10-1998
			JP	11507788 T	06-07-1999
			RU	2156035 C2	10-09-2000
			SE	9502142 A	14-12-1996
			WO	9642158 A1	27-12-1996
			US	5963564 A	05-10-1999
WO 9530318	A	09-11-1995	AU	706140 B2	10-06-1999
			AU	2301095 A	29-11-1995
			CA	2187876 A1	09-11-1995
			WO	9530318 A2	09-11-1995
			EP	1045557 A2	18-10-2000
			EP	0757881 A1	12-02-1997
			JP	9512404 T	09-12-1997
			SE	9603528 A	28-10-1996
			US	6269081 B1	31-07-2001
			US	5999528 A	07-12-1999
US 4958341	A	18-09-1990	AU	599860 B2	26-07-1990
			AU	3233789 A	12-10-1989
			CA	1312133 A1	29-12-1992
			DE	68926795 D1	14-08-1996
			DE	68926795 T2	09-01-1997
			EP	0335562 A2	04-10-1989
			ES	2088891 T3	01-10-1996
			HK	179396 A	04-10-1996
			JP	2013042 A	17-01-1990
			JP	2112391 C	21-11-1996
			JP	8021950 B	04-03-1996
			KR	204203 B1	15-06-1999
			SG	64883 A1	25-05-1999
			US	4872160 A	03-10-1989

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82